arXiv:1603.07302v2 [physics.chem-ph] 24 Mar 2016

Optimization of an exchange-correlation density functional for water

Michelle Fritz,[1] Marivi Fernández-Serra,[2,3] and José M. Soler[4]

[1)] *Departamento de Física de la Materia Condensada, Universidad Autónoma de Madrid, E-28049 Madrid, Spain*

[2)] *Department of Physics and Astronomy, Stony Brook University, Stony Brook, New York 11794-3800, USA*

[3)] *Institute for Advanced Computational Science, Stony Brook University, Stony Brook, New York 11794-3800, USA*

[4)] *Departamento e Instituto de Física de la Materia Condensada, Universidad Autónoma de Madrid, E-28049 Madrid, Spain*

(Dated: 25 March 2016)

We describe a method, that we call data projection onto parameter space (DPPS), to optimize an energy functional of the electron density, so that it reproduces a dataset of experimental magnitudes. Our scheme, based on Bayes theorem, constrains the optimized functional not to depart unphysically from existing *ab initio* functionals. The resulting functional maximizes the probability of being the "correct" parametrization of a given functional form, in the sense of Bayes theory. The application of DPPS to water sheds new light on why density functional theory has performed rather poorly for liquid water, on what improvements are needed, and on the intrinsic limitations of the generalized gradient approximation to electron exchange and correlation. Finally, we present tests of our water-optimized functional, that we call vdW-DF-w, showing that it performs very well for a variety of condensed water systems.

I. INTRODUCTION

Liquid water is arguably the most important substance for life, as well as for an immense number of problems of huge scientific and technological importance[1,2]. At the same time, despite its molecular simplicity, it is a liquid of astonishing complexity, with tens of thermodynamic anomalies. Ultimately, this complexity steams from the coexistence of covalent, electrostatic, and dispersion interactions, that have very different magnitudes but also subtle and critical interrelations[3–5]. Thus, it not surprising that electron density functional theory (DFT), a universal and completely non-empirical method, has had a particularly hard time in describing all these interactions with the required accuracy[3]. In fact, at present, DFT simulations cannot match the success of empirical force fields in simulating a wide range of effects and anomalies[6,7]. However, empirical methods are not necessarily reliable outside the range where they have been fitted, like in deeply undercooled water, a sate in which it has been predicted to have a liquid-liquid transition that could explain many of its anomalies at higher temperatures[8–10]. Therefore, its accurate description by DFT remains a very important challenge to understand the intricate structure and properties of liquid water. In this work we aim at uncovering in detail the deficiencies of present functionals and at developing an optimized functional within the generalized gradient approximation (GGA).

The parameterization and optimization of complex models is a pervasive problem in many areas, and in particular in the development of interatomic potentials and functionals for molecular dynamics. Generally, it requires to choose largely arbitrary functional forms that depend on many parameters, followed by a lengthy and difficult trial and error optimization[11]. The balance between the number of parameters and the size of the fitted data sets involves difficult and subjective decisions that are nevertheless critical to the results. Here we describe a general and powerful optimization scheme, data projection onto parameter space (DPPS), and its application to the optimization of an exchange-correlation (xc) energy functional of the electron density.

II. DATA PROJECTION ONTO PARAMETER SPACE

Before addressing functional optimization, it is useful to consider the method from a broader perspective. DPPS tries to find the optimal parameters for a complex model that "predicts" (calculates) a scalar function E (say the potential energy) that depends on a large number of variables $\mathbf{R}_i$ (say the atomic positions) and parameters. To be specific, imagine that we want to fit a pairwise interatomic potential $V(R)$ (assuming a single chemical species). The first step would be to choose a number of radial interpolation mesh points R_α for the interatomic distance. The parameters of the model would then be the values $\epsilon_\alpha \equiv V(R_\alpha)$, from which we would interpolate $V(R) = \sum_\alpha \epsilon_\alpha p_\alpha(R)$, where $p_\alpha(R)$ are a suitable set of interpolation basis functions. They are unambiguosly determined by the interpolation scheme (e.g. cubic splines[12,13]) and by the interpolation points: thus $p_\alpha(R)$ is the result of interpolating a function $f(R)$ with $f(R_\beta) = \delta_{\alpha\beta}$. Thus, $p_\alpha(R_\beta) = \delta_{\alpha\beta}$ and, for a sufficiently fine mesh, $p_\alpha(R) \simeq \delta(R - R_\alpha)$.

The next step is to set up a large data set of known system geometries $\mathbf{R}_i^n$ (where i denotes atoms and n denotes systems or geometries) and associated energies E_{ref}^n, obtained from experiments or from higher level calculations.

The intended prediction of the model would then be

$$E^n_{ref} = E^n \equiv \sum_{i<j} V(R^n_{ij}) = \sum_\alpha \epsilon_\alpha \sum_{i<j} p_\alpha(R^n_{ij}) \equiv \boldsymbol{\epsilon} \bullet \boldsymbol{\rho}^n \tag{1}$$

where $\boldsymbol{\epsilon} \equiv \{\epsilon_\alpha\}$ and $\boldsymbol{\rho}^n \equiv \{\rho^n_\alpha\}$, are vectors in "parameter space", with $\rho^n_\alpha \equiv \sum_{i<j} p_\alpha(R^n_{ij})$ being proportional to the radial distribution function (i.e. the radial interatomic density) at the mesh distances.

Eq. (1) sets the projections E^n_{ref} of the vector $\boldsymbol{\epsilon}$ of unknown parameters onto the vectors $\boldsymbol{\rho}^n$ of known data. Depending on the relative numbers of parameters and data, N_{dat} and N_{par}, the later will form a subspace of the former or they will overdetermine them. If $N_{dat} > N_{par}$, the system (1) can be solved in the least squares minimization sense, the solution being $\epsilon_\alpha = \sum_\beta S^{-1}_{\alpha\beta} u_\beta$, where $S_{\alpha\beta} \equiv \sum_n \rho^n_\alpha \rho^n_\beta$ and $u_\beta \equiv \sum_n E^n_{ref} \rho^n_\beta$.

In practice, even if $N_{dat} >> N_{par}$, all the data may nearly lie in a subspace of the parameter space. As a result, certain combinations of parameters may be poorly determined, and very large parameter values will result if we require an exact fit of the data energies. This situation will be apparent by a nearly singular matrix $S_{\alpha\beta}$, the standard cure being to invert it by singular value decomposition, discarding the subspace with small eigenvalues. Another standard alternative is to add a regularization penalty to large parameter changes, minimizing

$$Z = \sum_{n=1}^{N_{dat}} \left(\frac{E^n - E^n_{ref}}{\Delta E_n} \right)^2 + \sum_{\alpha=1}^{N_{par}} \left(\frac{\epsilon_\alpha - \epsilon^0_\alpha}{\Delta \epsilon_\alpha} \right)^2 = \min \tag{2}$$

where ΔE_n are error estimates of E^n_{ref}, and ϵ^0_α, $\Delta\epsilon_\alpha$ are initial estimates of the parameter values and their uncertainties. In practice, using $\epsilon^0_\alpha = 0$ and constant values for ΔE_n and $\Delta\epsilon_\alpha$ is essentially equivalent to singular value decomposition, with an eigenvalue cutoff $\sim \Delta E/\Delta\epsilon$. However, Eq. (2) provides a smoother and more natural transition to the Bayesian approach described below.

If no initial estimate ϵ^0_α is available, one can instead impose smoothness of $V(R)$, with a penalty to the first or second derivatives. Such a penalty can be rationalized, as in the Gaussian regression method[14], as prior Bayesian information on the function smoothness.

III. FUNCTIONAL OPTIMIZATION

Consider now the optimization of an xc energy functional of the electron density $\rho(\mathbf{r})$, in the generalized gradient approximation (GGA): $E_{xc}[\rho(\mathbf{r})] = \int \rho(\mathbf{r})\ \epsilon_{xc}(k_F(\mathbf{r}), k_G(\mathbf{r}))\ d^3\mathbf{r}$, where $k_F = (3\pi^2\rho)^{1/3}$ and $k_G = |\nabla\rho|/\rho$. We use these two wavevectors as functional variables, rather than the conventional ρ and $s = k_G/2k_F$, because they have the same dimension and similar magnitude (see below) and therefore they provide a more "isotropic" parameter space. To parametrize $\epsilon_{xc}(k_F, k_G)$, we choose interpolation points $k_{F\alpha}$ and $k_{G\beta}$ and we expand $\epsilon_{xc}(k_F, k_G) = \sum_{\alpha\beta}\ \epsilon_{\alpha\beta}\ p_\alpha(k_F)\ p_\beta(k_G)$, so that $\epsilon_{\alpha\beta}$ are our functional parameters. Assuming again that we have a dataset of system geometries $\mathbf{R}^n_i$ and corresponding total energies E^n_{ref}, we start with initial values $\epsilon^0_{\alpha\beta} = \epsilon^0_{xc}(k_{F\alpha}, k_{G\beta})$, where $\epsilon^0_{xc}(k_F, k_G)$ is a reference functional[15]. We find the self-consistent electron densities $\rho_n(\mathbf{r})$ and the initial total and xc energies, E^{0n}_{tot} and E^{0n}_{xc}, for each system n. Then, to first order in $\delta\epsilon_{\alpha\beta}$,

$$\begin{aligned} \delta E_{tot} &\simeq \sum_{\alpha\beta} \left(\frac{\partial E_{tot}}{\partial \epsilon_{\alpha\beta}} + \int \frac{\delta E_{tot}}{\delta \rho(\mathbf{r})} \frac{\partial \rho(\mathbf{r})}{\partial \epsilon_{\alpha\beta}} d^3\mathbf{r} \right) \delta\epsilon_{\alpha\beta} \\ &= \sum_{\alpha\beta} \frac{\partial E_{xc}}{\partial \epsilon_{\alpha\beta}} \delta\epsilon_{\alpha\beta} \simeq \delta E_{xc} \end{aligned} \tag{3}$$

where we have used that $\delta E_{tot}/\delta\rho(\mathbf{r}) = 0$ and $\partial E_{tot}/\partial\epsilon_{\alpha\beta} = \partial E_{xc}/\partial\epsilon_{\alpha\beta}$. Thus, we impose that $E^n_{ref} - E^{0n}_{tot} = E^n_{xc} - E^{0n}_{xc}$, or $E^n_{refxc} = E^n_{xc}$, where

$$E^n_{refxc} \equiv E^n_{ref} - E^{0n}_{tot} + E^{0n}_{xc} \tag{4}$$

and

$$E^n_{xc} \equiv \sum_{\alpha,\beta} \epsilon_{\alpha\beta} \int d^3\mathbf{r}\ \rho_n(\mathbf{r}) p_\alpha(k_F(\mathbf{r})) p_\beta(k_G(\mathbf{r})) = \boldsymbol{\epsilon} \bullet \boldsymbol{\rho}^n \tag{5}$$

where $\boldsymbol{\epsilon} = \{\epsilon_{\alpha\beta}\}$, $\boldsymbol{\rho}^n = \{\rho^n_{\alpha\beta}\}$, and $\rho^n_{\alpha\beta}$ are the integrals in (5). They are the density of electrons in parameter space, and they are closely related to the functions $g_1(r_s)$ and $g_3(s)$ of Zupan *et al*[16].

Although now the $\epsilon_{\alpha\beta}$'s span a 2-D interpolation grid (or a higher-dimensional mesh in case of a more complicated functional form, like meta-GGAs), it is convenient for notational simplicity to order them as a vector, using a single index index α for each pair of values $(k_F, k_G)_\alpha$, Thus, it is clear that Eqs. (1) and (5) are entirely equivalent. Then, we can solve (2) to find new functional parameters ϵ_α (i.e. a new xc functional) and iterate to obtain selfconsistency between ϵ_α, E^n_{xc}, and $\rho(\mathbf{r})$.

For exchange-only functionals, a simple scaling law requires the GGA exchange energy density to have the form $\epsilon^{GGA}_x(k_F, k_G) = \epsilon^{LDA}_x(k_F) F_x(s)$, where $\epsilon^{LDA}_x(k_F) = -3k_F/4\pi$ is the energy density in the local density approximation (LDA) and $s \equiv k_G/2k_F$ is a reduced adimensional gradient. Thus, like in the pair-potential example, in this case we must optimize the single-variable function $F_x(s)$, usually called "enhancement factor".

In general, known experimental data refer to energy differences rather than total energies. Thus we may know accurately reaction and atomization energies, or energy differences between different solid phases. To use these data efficiently, it suffices to substitute E_n and ϵ_n by ΔE_n and $\Delta\epsilon_n$ in the objective function Z and in all the equations above, where Δ refers to the difference between the two systems or geometries. This procedure also allows to use a large variety of structural and thermodynamic data. Thus, to impose known equilibrium geometries, we consider two geometries with one of the atoms displaced

by $\pm\Delta R$. Since the force must be zero, $\Delta\epsilon_n = 0$. The same can be done to impose zero pressure and stress at the known equilibrium geometry of a solid.

Equally, known vibration frequencies ω_q can be imposed using frozen-phonon displacements. Alternatively, a larger set of random geometries and corresponding energies (possibly taken from a molecular dynamics simulation) can be used to optimize both the equilibrium geometry and the deformation energies, beyond the harmonic approximation.

Electronic data can be imposed through constrained DFT[17]. For dipole moments and polarizabilities, the Lagrange multiplier associated to the electronic constraint is simply an external electric field. Thus, these constraints can be set similarly to those of the equilibrium geometry and deformation frequencies.

IV. BAYESIAN OPTIMIZATION

Although the simple penalty term of Eq. (2) avoids large parameter changes, in regions poorly determined by the data, it contains very little of our *ab initio* knowledge (and uncertainty) of how the functional should be. Therefore, the resulting functional, though optimal to fit a restricted set of data, may be rather unplausible from a theoretical point of view, and unreliable to reproduce other data. An obvious solution would be to minimize the error in the calculated energies, Eq. (2), under a number of specific theoretical constraints, not just a general penalty term. Some of those constraints should be "strict", while others may be more "relaxed" (quantitative). Thus, we know for sure that $F_x(0) = 1$, but we are much less certain about high values of s. More generally, the problem is to encode efficiently the known theoretical information, either strict or ambiguous, and this is exactly the aim of Bayesian probability theory.

Bayes theorem can be succinctly stated as

$$\mathcal{P}(theory|facts) = C\ \mathcal{P}(facts|theory)\mathcal{P}(theory), \quad (6)$$

where $\mathcal{P}(theory|facts)$ is the probability (or likelihood) that a theory is true, given that some facts have been observed, $\mathcal{P}(facts|theory)$ is the probability that those facts would be observed if the theory was true, $\mathcal{P}(theory)$ is (our estimate of) the *a priori* probability that the theory is true, and C is a normalization constant. In our case, "theory" means a quantitative parametrization of a given functional form (e.g. a set of GGA values $\epsilon_{\alpha\beta}$), "facts" are a set of E^n_{ref} and E^n_{GGA} energies, and "true" means "optimal to reproduce the energies" (not only of our dataset, but of all possible systems of interest).

Assuming a gaussian probability distribution,

$$\mathcal{P}(facts|theory) = C_1 \exp\left(-\frac{1}{2}\sum_n \left(\frac{E^n_{GGA} - E^n_{ref}}{\Delta E_n}\right)^2\right) \quad (7)$$

where ΔE_n are the expected errors in the computed energies, due to causes not related to the xc functional (e.g. basis set incompleteness), and to the inability of the GGA functional form to reproduce the exact energies. Equally

$$\mathcal{P}(theory) = C_2 \exp\left(-\frac{1}{2}\sum_{\alpha\beta}(\epsilon_\alpha - \epsilon^0_\alpha)C^{-1}_{\alpha\beta}(\epsilon_\beta - \epsilon^0_\beta)\right) \quad (8)$$

where ϵ^0_α is the average of $\epsilon_{xc}(k_{F\alpha}, k_{G\alpha})$ among different *ab initio* GGA functionals:

$$\epsilon^0_\alpha = \sum_i w_i \epsilon^i_\alpha \quad (9)$$

where index i labels different functionals, w_i are normalized weights assigned to them, and $\epsilon^i_\alpha \equiv \epsilon^i_{xc}(k_{F\alpha}, k_{G\alpha})$. C^{-1} is the inverse of the covariance matrix

$$C_{\alpha\beta} = \sum_i w_i(\epsilon^i_\alpha - \epsilon^0_\alpha)(\epsilon^i_\beta - \epsilon^0_\beta). \quad (10)$$

Notice that, in Eq. (8), we are using the discrepancies between the different "*ab initio*" functionals as a measure of our uncertainty on its exact form. However, the term $\mathcal{P}(theory)$ in Eq. (6) is important mostly in regions of the functional domain that are poorly sampled by the data, or in which the theoretical constraints are strict. In practice, the errors ΔE_n in Eq. (7) can be used as a knob to balance our relative uncertainties on energy data, on the one hand, and functional parameters, on the other. Notice also that the penalty approach of Eq. (2) is equivalent to using a diagonal covariance matrix for the functional parameters in the Bayesian approach. However, a diagonal covariance will not prevent the functional from developing strong oscillations, since the values of $\epsilon(k_F, k_G)$ at different points are assumed to be uncorrelated. In fact, the true covariance arising from a wide set of functionals is far from diagonal. Thus, Fig. (1) shows the average and the covariance of a set of 15 parameterizations[15,18–30] of the GGA exchange enhancement factor $F_x(s)$. It can be seen that the different functionals are highly correlated in a rather complicated way, with a very non-diagonal covariance.

Maximizing $\mathcal{P}(theory|facts)$, i.e.

$$Z = \sum_{n=1}^{N_{dat}} \left(\frac{E^n_{GGA}(\boldsymbol{\epsilon}) - E^n_{refxc}}{\Delta E_n}\right)^2 + \sum_{\alpha\beta}(\epsilon_\alpha - \epsilon^0_\alpha)C^{-1}_{\alpha\beta}(\epsilon_\beta - \epsilon^0_\beta) = \min \quad (11)$$

with respect to the functional parameters ϵ_α leads to a linear system of equations

$$\sum_\beta A_{\alpha\beta}\ \epsilon_\beta = B_\alpha \quad (12)$$

$$A_{\alpha\beta} = \sum_n \frac{\rho^n_\alpha \rho^n_\beta}{\Delta E^2_n} + C^{-1}_{\alpha\beta} \quad (13)$$

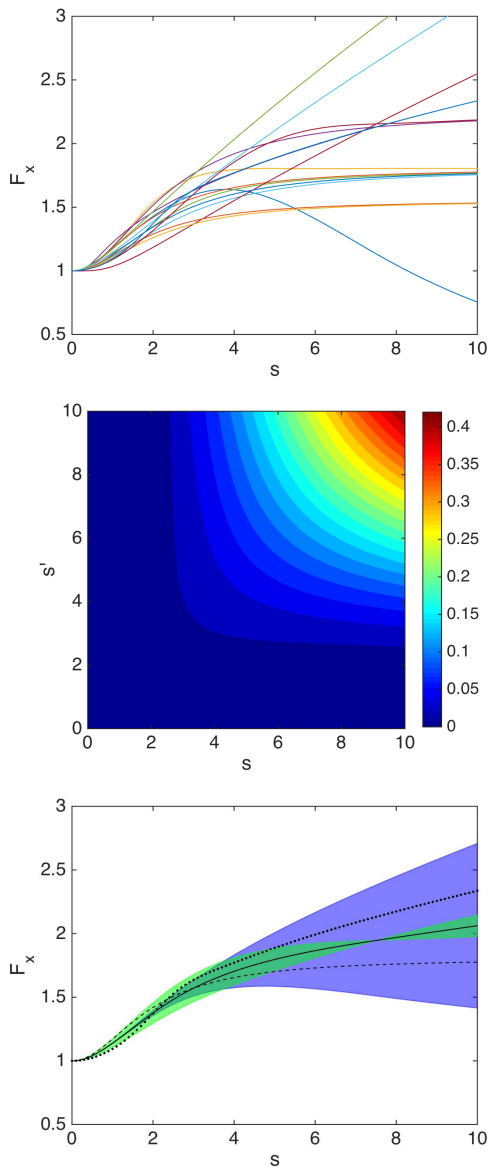

FIG. 1. Top: *Ab initio* GGA exchange enhancement factors $F_x(s)$ used as prior information for the Bayesian optimization.[15,18–30] Middle: Contour plot of the covariance matrix $\mathrm{Cov}[F_x(s), F_x(s')]$. Contour line values are multiples of 0.02. Bottom: Average (continuous line) and first two eigenvectors of the covariance matrix (multiplied by the square root of their corresponding eigenvalues and shown as half the width of the blue and green areas). Also shown are the enhancement factors of the PBE[15] (dashed line) and vdW-CX[30] (dotted line) functionals.

$$B_\alpha = \sum_n \frac{\rho_\alpha^n E_{refxc}^n}{\Delta E_n^2} + \sum_\beta C_{\alpha\beta}^{-1} \epsilon_\beta^0 \tag{14}$$

$$\rho_\alpha^n = \int d^3\mathbf{r}\ \rho_n(\mathbf{r})\ p_{F\alpha}(k_F(\mathbf{r}))\ p_{G\alpha}(k_G(\mathbf{r})) \tag{15}$$

The resulting Bayesian method for functional optimization shares many methodological characteristics of the BEEF method of Wellendorff *et al*[31–33], as well as with that of Aldegunde *et al*[34]. Like them, it has the ability to quantify the errors due to the uncertainty of the resulting functional, although we will not discuss this ability in the present work. However, the prior information in those methods is "objective", i.e. of general mathematical character, like a requirement of smoothness for the resulting functional. In contrast, our method uses an "informative" prior, i.e. functional constraints and uncertainties arising from a variety of *ab initio* criteria.

V. DATA ANALYSIS IN PARAMETER SPACE

An important component of our approach is projecting the electron density $\rho(\mathbf{r})$ of the data (i.e. of the water molecules) into "parameter space" which, for a GGA functional, is that spanned by wavevectors k_F and k_G. For a given geometry, this is

$$\rho(k_F, k_G) = \int d^3\mathbf{r}\rho(\mathbf{r})\delta\left((3\pi^2\rho(\mathbf{r}))^{1/3} - k_F\right) \times\ \delta\left(\frac{|\nabla\rho(\mathbf{r})|}{\rho(\mathbf{r})} - k_G\right). \tag{16}$$

Figure 2 shows $\rho(k_F, k_G)$ for the isolated water molecule in its equilibrium geometry. It may be seen that the

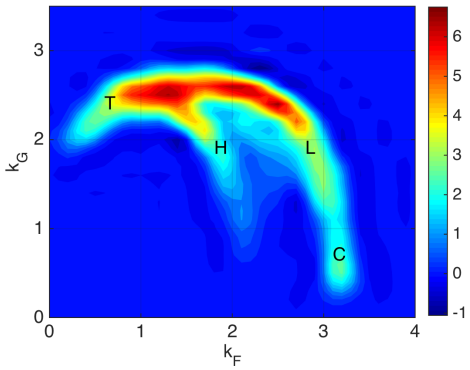

FIG. 2. Histogram of the electron density of the water molecule, in the equilibrium geometry, as a function of wavevectors $k_F = (3\pi^2\rho)^{1/3}$ and $k_G = |\nabla\rho|/\rho$. The labels C, H, L, and T indicate regions of parameter space dominated by the real-space regions of the oxygen core, hydrogen atoms, lone pairs, and electron tails, respectively. Atomic units are used.

most relevant region of this parameter space is that with $(k_F^2 + k_G^2)^{1/2} \sim$ const. We have observed that this is also true for a variety of systems, not only water. However, what is generally most important are the *changes* of electron density between different systems and geometries. Thus, Fig. 3 shows the change in $\rho(k_F, k_G)$ between the molecule in its equilibrium geometry with the vdW-DF-cx functional[30] and with the experimental geometry. The basic difference between these two geometries

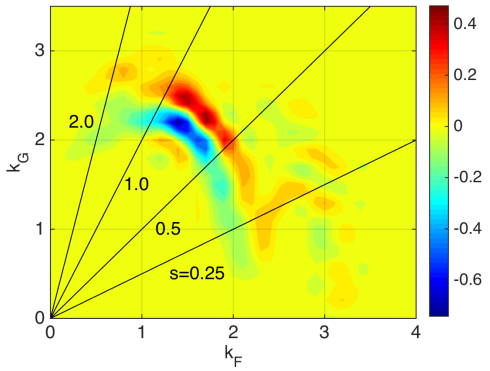

FIG. 3. Difference between the electron density $\rho(k_F, k_G)$ of the water molecule in the experimental geometry minus that in the relaxed geometry, using the vdW-DF-cx functional in both cases. Also shown are lines of constant reduced gradient $s = k_G/2k_F$.

is that, with vdW-DF-cx, the OH bonds are $\sim 1.3\%$ too long. This translates into a "radial" shift in $\rho(k_F, k_G)$, nearly insensitive to the only variable s of GGA exchange. This implies that optimizing $F_x(s)$ to reproduce the correct monomer geometry will force the functional to use rather artificial mechanisms, based on subtle changes of $\rho(k_F, k_G)$, to achieve the bond contraction.

Fig. 4 shows the changes in $\rho(k_F, k_G)$ upon formation of a water dimer. In this case, the differences are more dependent on s, and therefore more susceptible to the optimization of $F_x(s)$. Equally, Fig. 5 shows the deformation density upon formation of a water trimer:

$$\Delta\rho_{123} = \rho_{123} - \rho_{12} - \rho_{23} - \rho_{13} + \rho_1 + \rho_2 + \rho_3 \qquad (17)$$

where ρ_i, ρ_{ij}, and ρ_{123} are the densities of the monomers, dimers, and trimer, respectively. Comparison between Figs. 4 and 5 shows that the dimer and trimer deformation densities are, to a large extent, opposite, implying that two- and three-body xc energies will have a large compensation.

VI. OPTIMIZATION OF A FUNCTIONAL FOR WATER

Using the DPPS methodology, we have optimized a GGA exchange functional for water. Of course, given the narrow scope of our system, we cannot assume that the resulting functional is significant for a broader range of DFT applications. Rather, our aim is to shed new light on why DFT has been so frustratingly poor in simulating liquid water, on what changes are needed to improve it, and on the intrinsic shortcomings of the GGA form for this important system.

For the DFT calculations, we have used a DPPS-enhanced version of the SIESTA code[35], with a quadruple-zeta plus double polarization basis set[36] and a highly converged mesh for real-space integrals and Fourier transforms.

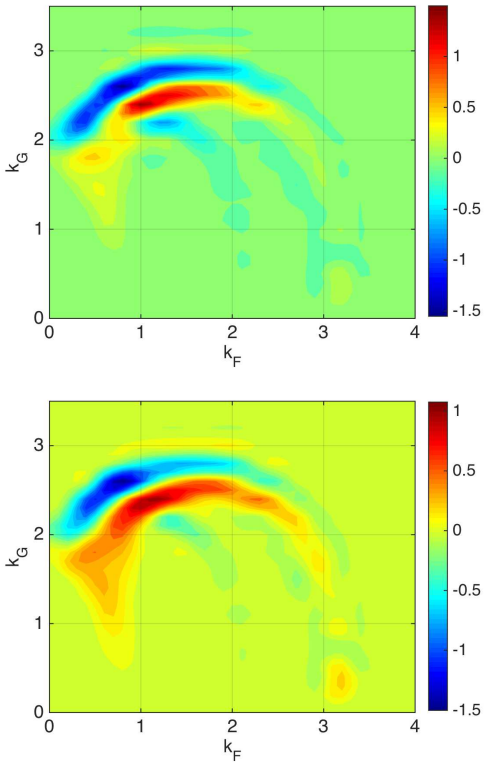

FIG. 4. Difference $\Delta\rho_{12} = \rho_{12} - \rho_1 - \rho_2$ between the electron density $\rho(k_F, k_G)$ of a water dimer in its equilibrium geometry (top) and in a non-hydrogen-bonded geometry (bottom), and that of the two separated monomers (in their dimer geometry).

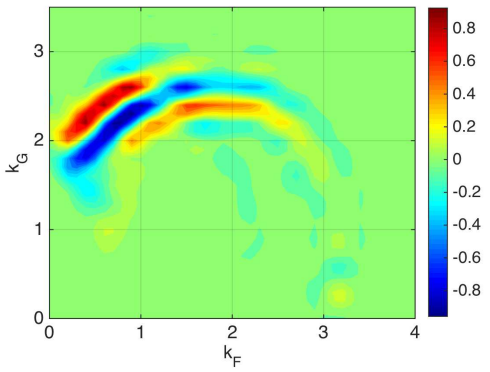

FIG. 5. Deformation density, defined by Eq. (17), upon formation of a water trimer in its equilibrium geometry.

As prior information for the Bayesian optimization, we have used the 15 enhancement factors shown in Fig. 1. For the reference energies, we use the MB-pol force field[11,37–39], a sophisticated fit to highly accurate quantum chemical calculations of water monomers and clusters. We selected geometries of 300 monomers, 330 dimers, and 100 trimers for our training set. The monomers were generated with regular grids of bond dis-

tances and angles. The dimers and trimers were taken from the MB-pol training set (see SI for more information). It is important to emphasize that, although fitted only to monomers, dimers and trimers, the MB-pol force field has been shown to reproduce accurately the structural and thermodynamic properties of condensed phases of water[11,37–39]. This shows that two- and three-body interactions are dominant in water, and it gives us confidence that fitting (a large subset of) the MB-pol training set will also produce an accurate exchange functional for water. Details of these geometries and of our fitting weights and procedures can be found in the supporting information (SI)[40].

Among the 15 functionals used as priors, some were designed for a limited set of systems (e.g. bulk solids). Furthermore, while some are thoroughly tested and truly *ab initio*, others are more marginal, or they contain some degree of empirical information. Therefore, rather than the unweighted mean of the 15 functionals, we decided to use the newly proposed vdW-DF-cx[41] as the reference functional (ϵ_α^0 parameters in Eqs. (2) and (8)). This functional keeps the non local correlation form of the original vdW-DF[42] and it proposes a new consistent exchange (cx) which has shown to improve over the previous versions of vdW-DF.

Figure 6 shows the exchange enhancement functions, optimized for water monomers, with and without Bayesian constraints (Eqs. (11) and (2), respectively): Also shown is the enhancement factor of the vdW-DF-cx

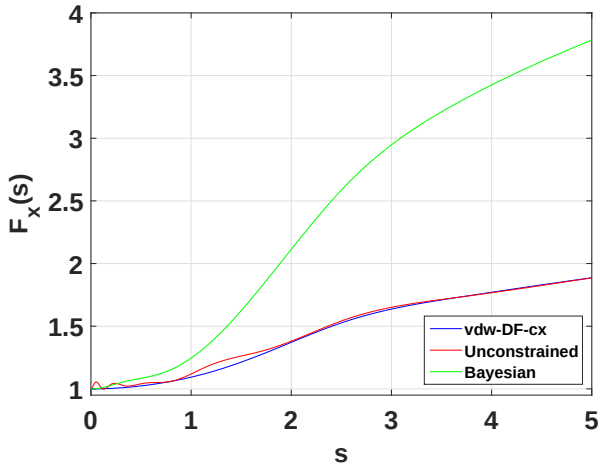

FIG. 6. GGA exchange enhancement factor $F_x(s)$, optimized to reproduce reference energies[43] of the water monomer at 300 geometries (red). We show the result with (green) and without (red) bayesian constrains. The enhancement factor of the vdW-DF-cx functional (blue), used as the reference functional in the optimization, is also shown for comparison.

functional[30], used as reference for the initial functional parameters (values ϵ_α^0 in Eqs. (2) and (8)). Although the resulting energies are in excellent agreement with the reference data, the unconstrained functional shows strong oscillations, that make it rather unplausible from a theoretical point of view. This is of course a consequence of not having imposed any physical nor mathematical constraints on the plausible shapes of $F_x(s)$. When we do impose such constraints, by performing a Bayesian optimization (Eq. (11)) with *ab initio* functions $F_x(s)$ as prior information, the result is much more plausible, but there is a dramatic increase of $F_x(s)$. To understand this result, we notice in Fig. 3 that the correct density $\rho(k_F, k_G)$ (corresponding to shorter OH distances) is shifted to larger values of k_G (and therefore, to larger s) for $s \sim 1$, in the region dominated by the electron tails. Therefore, a shorter bond is favoured by a larger slope of $F_x(s)$ for $s \sim 1$. And, since $F_x(s)$ is strongly constrained to be monotonous for $s < 4$, the result is a large overall increase.

The previous argument is in contrast with the usual result, that GGAs increase binding distances relative to the LDA. This occurs because longer bonds decrease k_F (the local density) and increase k_G (its gradient) in the bond region (for an example, see the case of the oxygen molecule in SI). This increases s, what is favoured by the positive slope of $F_x(s)$ in the GGAs. The departure of H_2O from this usual behavior is due to the absence of a hydrogen core and to the small OH distance, with the hydrogen atom practically buried within the oxygen density. The main change, upon decreasing the OH distance, is a global electronic contraction, with an average *increase* of the density gradient around $s \sim 1$ (Fig. 3). Therefore, the positive slope of $F_x(s)$ favors a shorter OH distance, which *contracts* by 0.2% with GGA-PBE, relative to the LDA.

Although the Bayesian functional of Fig. 6 is not itself physically unplausible, its shape is very different from that required to fit the interaction energies of dimers and trimers (Fig. 7). As a consequence, we have found essentially impossible to fit simultaneously the monomer geometry and the interaction energies. This has led us to introduce an *ad hoc* correction for the energy difference between the exact and GGA energies of each monomer. This correction has the functional form of Ref. 43 and it adds a negligible overhead to the DFT calculation. Similar monomer corrections have been previously applied by other authors[44,45]. Results assesing the qualty of this 1-body correction are presented in the SI[40]. Although rather unsatisfactory from a theoretical point of view, this approach allows us to proceed with the optimization of the functional, in order to reproduce the more relevant energies of interaction between the different water molecules. In the future, we expect richer functional forms, like those of meta-GGAs, to correct this deficiency.

Figure 7 shows $F_x(s)$ optimized to reproduce the interaction energies of water dimers and trimers, using the Bayesian method described above. The importance of accurate 3-body term corrections to GGA functionals for water has already been highlighted in Ref. 44. The resulting functional also shows a dramatic improvement in the fitted energies, while the shape of $F_x(s)$ is still within the bounds of physical and mathematical plausibility. In

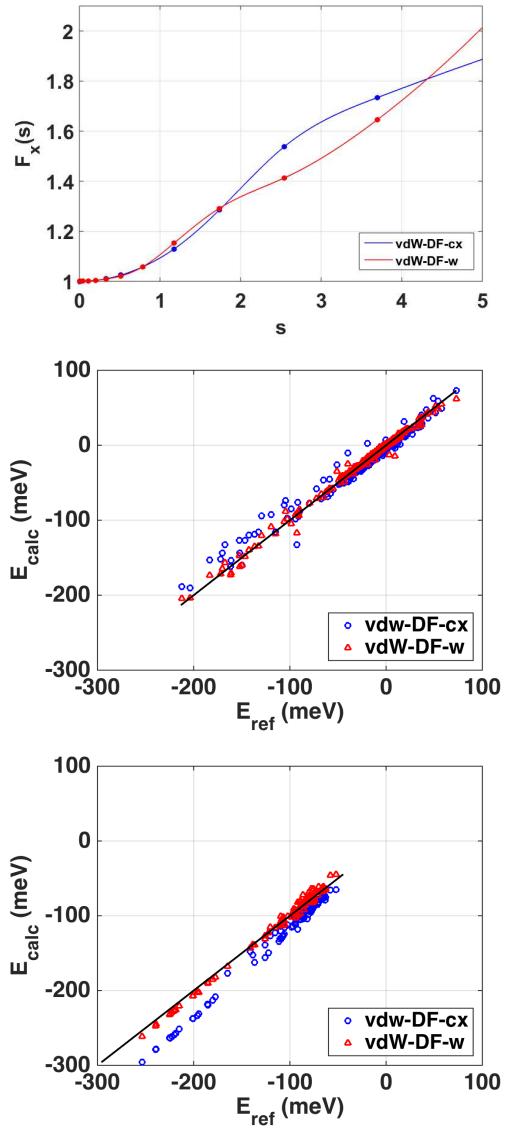

FIG. 7. Top: GGA exchange enhancement factor, optimized to reproduce reference energies of 330 water dimers and 100 trimers (red). Enhancement factor of the vdW-DF-cx functional (blue). Middle, red dots: comparison between the dimer interaction energies $\Delta E_{12} = E_{12} - E_1 - E_2$, calculated with the optimized functional, and the corresponding reference energies. The black line corresponds to a perfect agreement. The mean and mean square errors are $\langle \Delta E \rangle = 0.02$ meV, $\langle \Delta E^2 \rangle^{1/2} = 0.62$ meV. The blue dots are the same for the vdW-DF-cx functional, with $\langle \Delta E \rangle = -0.17$ meV, $\langle \Delta E^2 \rangle^{1/2} = 1.33$ meV. Bottom: the same for the trimer interaction energies $\Delta E_{123} = E_{123} - E_{12} - E_{13} - E_{23} + E_1 + E_2 + E_3$. For the optimized functional, $\langle \Delta E \rangle = 0.00$ meV, $\langle \Delta E^2 \rangle^{1/2} = 0.60$ meV. For the vdW-DF-cx functional, $\langle \Delta E \rangle = -1.85$ meV, $\langle \Delta E^2 \rangle^{1/2} = 2.21$ meV.

the following, we will denote this functional vdW-DF-w, or van der Waals density functional optimized for water.

In addition to the energies of monomers, dimers, and trimers, the MB-pol force field was parameterized to reproduce also the electric dipole μ and isotropic polarizability α of the molecule, which are important for the long range interactions. Therefore, it is important to notice that, although our vdW-DF-w functional was not trained to reproduce μ and α, it nevertheless does so rather well. Thus we find $\mu = 1.81$ D, $\alpha = 1.41$ Å^3 for vdW-DF-cx and $\mu = 1.81$ D, $\alpha = 1.39$ Å^3 for vdW-DF-w, versus the reference values[37,43] of $\mu = 1.86$ D and $\alpha = 1.43$ Å^3.

VII. RESULTS

In the following section we will present results for gas and condensed phases of water, obtained with our optimized vdW-DF-w functional, trained to reproduce 2- and 3-body energies from our training set. For comparison, we will also present results for the vdW-DF-cx functional[41], which has not yet been sufficiently evaluated for water.

As a first test of the ability of vdW-DF-w to describe water systems beyond trimers, we show in Fig. 8 the energies of various isomers of the water hexamer (whose geometries are shown in SI). We can see that our functional does improve considerably the energies of the vdW-DF-cx functional, relative to the MB-pol reference energies (that reproduce accurately those of high-level quantum chemistry calculations). Thus, in contrast with vdW-DF-cx, our functional reproduces the correct order of the isomer binding energies.

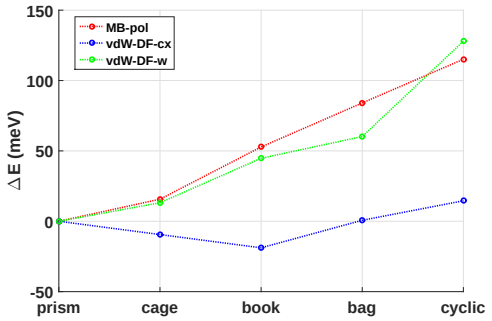

FIG. 8. Relative binding energies of five isomers of the H2O hexamer, whose geometries are shown in the SI. We show the energies for vdW-DF-cx, for our optimized vdW-DF-w functional, and for the reference MB-pol force field.

Because the small energy differences between its many solid structures, the phase diagram of water ice is a major challenge for DFT. We have computed lattice energies for different ice polymorphs at fixed geometries. For the vdW-DF and MB-pol calculations, the geometries are those given by Santra *et al*[46] for the PBE0-vdW-DF^{TS} functional. For PBE, they are those given in the same reference for this functional. These geometries (given in SI) are rather different from those of MB-pol, which results in an incorrect ordering of the different polymorphs. It

is important to emphasize that, although no results have been reported for relaxed MB-pol ice energies, a thorough comparison between MB-pol and Quantum Monte Carlo energies for liquid snapshots demonstrated that this force field yields much better agreement than any density functional[39]. In Table I we compare our results with those of MB-pol and of various other functionals. It is significant that vdW-DF-w is the only functional that reproduces qualitatively the behaviour of our reference MB-pol energies.

TABLE I. Lattice energies of different ice phases, relative to ice Ih, in meV/molecule. All the calculations were done for fixed lattice and molecular geometries (provided in the SI). Δ_1 indicates that the monomer correction described in the text was included.

	Ih	II	IX	XIII	XIV	XV	VIII
vdW-DF-w	0	-11	-14	-22	-23	-7	11
vdW-DF-w-Δ_1	0	-39	-17	-10	-13	-18	-5
vdW-DF-cx	0	30	14	20	24	49	90
vdW-DF-cx-Δ_1	0	22	11	31	34	38	73
vdW-DF1	0	12	8	13	14	23	50
vdW-DF2	0	4	6	5	5	14	33
PBE	0	72	51	83	96	114	180
PBE[46]	0	69	49	80	93	110	177
MB-POL	0		-37		-48	-91	-77

It is well established that the competition between local structures, with different average densities (generally called low- and high-density liquids, LDL and HDL), is the origin of abnormally high response functions and of most water anomalies[9,10]. This competition depends critically on a very subtle balance between the hydrogen-bond and van der Waals interactions, and it represents an extremely difficult challenge for liquid water simulations in general, and for *ab initio* molecular dynamics (AIMD) in particular. We have performed AIMD of 128 water molecules at normal temperature and pressure conditions, using SIESTA with a double-zeta polarized basis set. Starting from the last geometry of a long Tip4P-2005[47] simulation, the system was thermalized with AIMD for 5 ps, and then sampled for 10 ps. This procedure was followed for vdW-DF-cx and for our optimized vdw-DF-w functional, with and without correcting for the monomer energies. Nuclear quantum effects were not included. The resulting oxygen-oxygen radial distribution functions are compared in Fig. 9 with experiment and with an MB-pol simulation without nuclear quantum effects. The later was performed under identical conditions and parameters as the AIMD simulations. It can be seen that our vdW-DF-w functional compares much better than vdW-DF-cx with the reference MB-pol simulation. The optimized functional overcorrects the underestimation of O-O distances (first peak of the RDF) by the vdW-DF-cx functional, giving a right-shifted peak.

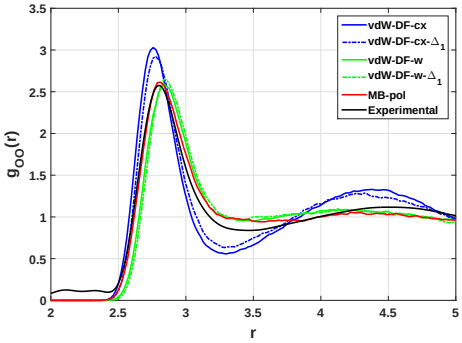

FIG. 9. Oxygen-oxygen radial distribution function of water, at normal temperature-pressure conditions, with the vdW-DF-cx and optimized vdW-DF-w functionals, as well as with the reference MB-pol force field. Notice that our MB-pol simulation was performed under the same conditions as the DFT ones, and it is much shorter than that of Ref. 11. The experimental RDF is also shown for comparison[48]. Nuclear quantum effects were not included in any of the simulations.

The correction of the potential energy surface of the monomer dominates the total energies and it reduces the intramolecular OH distance by 1.3%. From the well known anticorrelation between intra- and inter-molecular bond distances and energies, it may be expected that the correction weakens the hydrogen bonds, and favors the HDL configurations. However, it can be seen that in practice the monomer correction has a rather small effect on the RDF.

Perhaps the most blatant shortcoming of DFT simulations of liquid water is the huge range of discrepancies in the equilibrium density, ranging from $\sim$ 0.7 g/cm^3 for some GGAs to $\sim$ 1.15 g/cm^3 for some vdW-DFs[3]. The pressure-density curves from our simulations are compared with experiment in Fig. 10. It may be seen

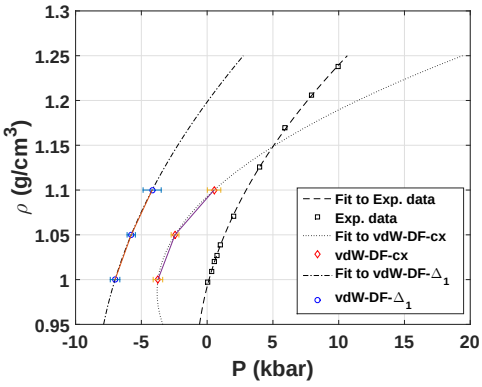

FIG. 10. Comparison of the density-pressure curves from our liquid water simulations with vdW-DF-cx and vdW-DF-w and from experiment[49].

that vdW-DF-cx and vdW-DF-w overestimate the equi-

librium density by 10% and 20%, respectively. In the latter, however, this is because of a rigid shift of the curve to lower pressures, so that the compressibility (shown in SI) is in good agreement with experiment. Since the pressure is also very sensitive to finite size and nuclear quantum effects (not included in our simulations), as well as to basis set superposition errors (the DZP basis used in the simulations is smaller than that used in the functional optimization), it remains to be studied to what extent the density overestimation is due to these effects.

Another important magnitude, that depends strongly on the ratio of LDL to HDL in the actual liquid, is the self diffusion constant. The experimental value of 0.23 Å^2/ps is strongly underestimated by GGAs under normal conditions of density and temperature[3]. For the vdW-DF-cx functional, we find 0.08 Å^2/ps. When we include the monomer energy correction, this increases slightly to 0.09 Å^2/ps, still much lower than the experimental value. With vdW-DF-w, we obtain 0.17 Å^2/ps without the monomer correction and 0.23 Å^2/ps with it, in agreement with experiment.

VIII. CONCLUSIONS

In conclusion, we have designed and implemented a general Bayesian method to optimize an exchange-correlation functional. It combines *a priori* information from *ab initio* functional constraints, with a database of reference energies and geometries. Using this method, we have optimized for water a GGA exchange functional, combined with the correlation functional of vdW-DF[42]. As prior information, we have used the covariance of 15 GGA exchange functionals. As reference, we used a database of hundreds of monomers, dimers, and trimers taken from the MB-pol training set[11,37–39]. We find that the optimization of the monomer geometry is largely incompatible with that of intermolecular interactions. As an *ad hoc* solution, we have developed a correction for the potential energy surface of the monomer, and we have optimized the functional to reproduce the dimer and trimer interaction energies. The resulting functional performs considerably better than previous GGA and vdW-DF functionals for extended clusters, ices, and liquid water.

We acknowledge Prof. M. Gillan for discussions and for providing an initial database of MP2 energies and geometries. We also acknowledge Prof. F. Paesani for the MB-pol code and for assistance in its installation and use. This work has been funded by MINECO grants FIS2012-37549 and FIS2015-64886 and by DOE grants DE-FG02-09ER16052 and de-sc0003871.

[1] G. Hummer and A. Tokmakoff, The Journal of Chemical Physics **141**, 22 (2014).
[2] E. M. J. B. S. M. Q. S. E. Walter, M.G. Warren and N. Lewis, Chem. Rev. **110**, 6446 (2010).
[3] J. Wang, R. Pérez, J. M. Soler, E. Artacho, and M.-V. Fernández-Serra, J. Chem. Phys. **134**, 024516 (2011).
[4] B. Pamuk, J. M. Soler, R. Ramírez, C. P. Herrero, P. W. Stephens, P. B. Allen, and M.-V. Fernández-Serra, Phys. Rev. Lett. **108**, 193003 (2012), URL http://link.aps.org/doi/10.1103/PhysRevLett.108.193003.
[5] F. Corsetti, E. Artacho, J. M. Soler, S. S. Alexandre, and M.-V. Fernndez-Serra, The Journal of Chemical Physics **139**, 194502 (2013).
[6] C. Vega, J. L. F. Abascal, M. M. Conde, and J. L. Aragones, Faraday Discuss. **141**, 251 (2009), URL http://dx.doi.org/10.1039/B805531A.
[7] C. Vega and J. L. F. Abascal, Phys. Chem. Chem. Phys. **13**, 19663 (2011), URL http://dx.doi.org/10.1039/C1CP22168J.
[8] F. E. U. Poole, P. H. Sciortino and S. H. E., Nature **360**, 324 (1992).
[9] L. Xu, P. Kumar, S. V. Buldyrev, S.-H. Chen, P. H. Poole, F. Sciortino, and H. E. Stanley, Proceedings of the National Academy of Sciences of the United States of America **102**, 16558 (2005), http://www.pnas.org/content/102/46/16558.full.pdf, URL http://www.pnas.org/content/102/46/16558.abstract.
[10] J. L. F. Abascal and C. Vega, The Journal of Chemical Physics **133**, 234502 (2010).
[11] G. R. Medders, V. Babin, and F. Paesani, J. Chem. Theor. Comput. **10**, 2906 (2014).
[12] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes* (Cambridge University Press, Cambridge, 1992).
[13] In the cubic splines method, it is also necessary to specify the two boundary conditions (either a given value of the first derivative, usually zero, or a zero second derivative.
[14] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning* (MIT Press, Cambridge, MA, 2006).
[15] J. P. Perdew, K. Burke, and M. Ernzerhof, Phys. Rev. Lett. **77**, 3865 (1996).
[16] A. Zupan, K. Burke, M. Ernzerhof, and J. P. Perdew, J. Chem. Phys. **106**, 10184 (1997).
[17] B. Kaduk, T. Kowalczyk, and T. V. Voorhis, Chem. Rev. **112**, 321 (2012).
[18] J. P. Perdew and W. Yue, Phys. Rev. B **33**, 8800 (1986), **33**, 8800(R) (1986).
[19] A. D. Becke, Phys. Rev. A **38**, 3098 (1988).
[20] Y. Zhang and W. Yang, Phys. Rev. Lett. **80**, 890 (1998).
[21] B. Hammer, L. B. Hansen, and J. K. Nørskov, Phys. Rev. B **59**, 7413 (1999), URL http://link.aps.org/doi/10.1103/PhysRevB.59.7413.
[22] Z. Wu and R. E. Cohen, Phys. Rev. B **73**, 235116 (2006), URL http://link.aps.org/doi/10.1103/PhysRevB.73.235116.
[23] J. P. Perdew, A. Ruzsinszky, G. I. Csonka, O. A. Vydrov, G. E. Scuseria, L. A. Constantin, X. Zhou, and K. Burke, Phys. Rev. Lett. **100**, 136406 (2008).
[24] L. S. Pedroza, A. J. R. da Silva, and K. Capelle, Phys. Rev. B **79**, 201106 (2009).
[25] M. M. Odashima, K. Capelle, and S. B. Trickey, J. Chem. Theor. Comp. **5**, 798 (2009).
[26] A. E. Mattsson and R. Armiento, Phys. Rev. B **79**, 155101 (2009), URL http://link.aps.org/doi/10.1103/PhysRevB.79.155101.
[27] E. D. Murray, K. Lee, and D. C. Langreth, J. Chem. Theor. Comp. **5**, 2754 (2009).
[28] J. Klimes, D. R. Bowler, and A. Michaelides, J. Phys.: Condens. Matter **22**, 022201 (2010).
[29] V. R. Cooper, Phys. Rev. B **81**, 161104 (2010), URL http://link.aps.org/doi/10.1103/PhysRevB.81.161104.
[30] K. Berland and P. Hyldgaard, Phys. Rev. B **89**, 035412 (2014), URL http://link.aps.org/doi/10.1103/PhysRevB.89.035412.
[31] J. Wellendorff, K. T. Lundgaard, A. Mogelhoj, V. Petzold, D. D. Landis, J. K. Norskov, T. Bligaard, and K. W. Jacobsen, Phys. Rev. B **85**, 235149 (2012).
[32] J. Wellendorff, K. T. Lundgaard, K. W. Jacobsen, and T. Bligaard, J. Chem. Phys. **140**, 144107 (2014).
[33] A. J. Medford, J. Wellendorff, A. Vojvodic, F. Studt, F. Abild-Pedersen, K. W. Jacobsen, T. Bligaard, and J. K. Norskov, Science **345**, 6193 (2014).

[34] M. Aldegunde, J. R. Kermode, and N. Zabaras, J. Comp. Phys. **311**, 173 (2016).
[35] J. M. Soler, E. Artacho, J. D. Gale, A. García, J. Junquera, P. Ordejón, and D. Sánchez-Portal, J. Phys.: Condens. Matter **14**, 2745 (2002).
[36] F. Corsetti, M.-V. Fernandez-Serra, J. M. Soler, and E. Artacho, J. Phys. Condens. Matter **25**, 435504 (2013).
[37] V. Babin, C. Leforestier, and F. Paesani, J. Chem. Theor. Comput. **9**, 5395 (2013).
[38] V. Babin, G. R. Medders, and F. Paesani, J. Chem. Theor. Comput. **10**, 1599 (2014).
[39] G. R. Medders, A. W. Götz, M. Morales, P. Bajaj, and F. Paesani, J. Chem. Phys. **143**, 104102 (2015).
[40] Supporting information of this article.
[41] K. Berland and P. Hyldgaard, Phys. Rev. B **89**, 035412 (2014).
[42] M. Dion, H. Rydberg, E. Schröder, D. C. Langreth, and B. I. Lundqvist, Phys. Rev. Lett. **92**, 246401 (2004).
[43] H. Patridge and D. W. Schwenke, J. Chem. Phys. **106**, 4618 (1997).
[44] M. J. Gillan, D. Alfè, A. P. Bartók, and G. Csányi, J. Chem. Phys. **139**, 244504 (2013).
[45] A. P. Bartók, M. J. Gillan, F. R. Manby, and G. Csányi, Phys. Rev. B **88**, 054104 (2013).
[46] B. Santra, J. Klimes, A. Tkatchenko, D. Alfé, B. Slater, A. Michaelides, R. Car, and M. Scheffler, J. Chem. Phys. **139**, 154702 (2013).
[47] J. L. F. Abascal and C. Vega, J. Chem. Phys. **123**, 234505 (2005).
[48] L. B. Skinner, C. Huang, D. Schlesinger, L. G. M. Pettersson, A. Nilsson, and C. J. Benmore, The Journal of Chemical Physics **138**, 074506 (2013).
[49] W. Wagner and A. Pruss, J. Phys. Chem. Ref. Data **31**, 387 (2002).